arXiv:1405.2566v1 [stat.ML] 11 May 2014

Learning Modular Structures from Network Data and Node Variables

Elham Azizi ELHAM@BU.EDU
James E. Galagan JGALAG@BU.EDU
Departments of Biomedical Engineering and Microbiology
Boston University
Boston, MA 02215, USA

Edoardo M. Airoldi AIROLDI@FAS.HARVARD.EDU
Department of Statistics
Harvard University
Cambridge, MA 02138, USA

Abstract

A standard technique for understanding underlying dependency structures among a set of variables posits a shared conditional probability distribution for the variables measured on individuals within a group. This approach is often referred to as module networks, where individuals are represented by nodes in a network, groups are termed modules, and the focus is on estimating the network structure among modules. However, estimation solely from node-specific variables can lead to spurious dependencies, and unverifiable structural assumptions are often used for regularization. Here, we propose an extended model that leverages direct observations about the network in addition to node-specific variables. By integrating complementary data types, we avoid the need for structural assumptions. We illustrate theoretical and practical significance of the model and develop a reversible-jump MCMC learning procedure for learning modules and model parameters. We demonstrate the method accuracy in predicting modular structures from synthetic data and capability to learn influence structures in twitter data and regulatory modules in the *Mycobacterium tuberculosis* gene regulatory network.

Keywords: Module Networks, Blockmodels, Gene Regulatory Networks, ChIP-Seq, Reversible-Jump MCMC, Data Integration

1. Introduction

There is considerable interest in modeling dependency structures in a variety of applications. Examples include reconstructing regulatory relationships from gene expression data in gene networks or identifying influence structures from activity patterns such as purchases, posts, tweets, etc in social networks. Common approaches for learning dependencies include using Bayesian networks and factor analysis (Koller and Friedman, 2009).

Module networks (Segal et al., 2005, 2003) have been widely used to find structures (e.g. gene regulation) between groups of nodes denoted as modules, based on measurements of node-specific variables in a network (e.g. gene expression). The motivation lies in that nodes that are influenced or regulated by the same parent node(s), have the same conditional probabilities for their variables. For example, in gene regulatory networks, groups of genes respond in concert under certain environmental conditions (Qi and Ge, 2006) and are thus likely to be regulated by the same mechanism. In

other domains, such as social networks, communities with similar interests or affiliations may have similar activity in communicating messages in response to news-outbreaks or similar purchases in response to marketing advertisements (Kozinets, 1999; Aral et al., 2009).

However, inferring dependencies merely from node-specific variables can lead to higher rate of false positives (Michoel et al., 2007). For example, a dependency might be inferred between two unrelated nodes due to existing confounding variables. This can introduce arbitrary or too many parents for a module. To avoid over-fitting in inferring module networks, additional structural assumptions such as setting the maximum number of modules or maximum number of parents per module may be required. This in turn presents additional inductive bias and results become sensitive to assumptions. Moreover, searching through the entire set of candidate parents for each module is computationally infeasible.

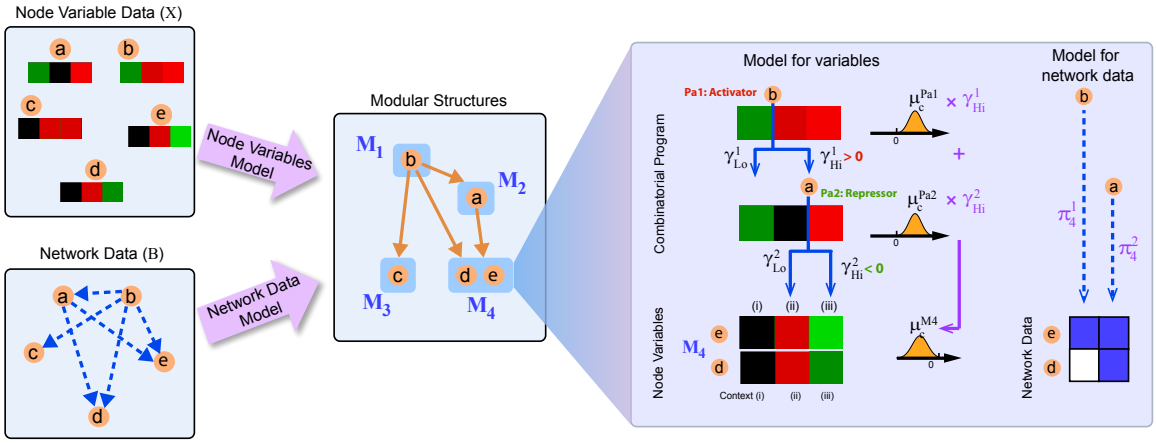

Figure 1: **Illustration of proposed model:** Modular structures are learned from node variables (e.g. gene expression) and network data (e.g. protein-DNA interactions). Node variables are color-coded ranging from green (low) to red (high). A number of parents are assigned to each module (orange links). A combinatorial program is inferred for each module; example shown for module M_4.

Alternatively, we can take advantage of existing network data and by integrating node interactions with node variables, we can avoid structural assumptions. For example, to learn gene regulatory networks, we can use protein-DNA interaction data, which shows physical interactions between proteins of genes (known as Transcription Factors) with promoter regions of other genes, leading to regulation of transcription (and expression) of the latter genes. This data can be measured using chromatin immunoprecipitation of DNA-bound proteins, i.e. ChIP-ChIP or ChIP-Seq technologies, which have shown to be informative of regulation (Galagan et al., 2013; Liu et al., 2013; Celniker et al., 2009; Yeang et al., 2004). As another example, to learn influence structures in a twitter network, we can integrate the network of who-follows-who with measurements of users activities.

Identifying modules or block structures from network data has been well-studied, e.g., using stochastic blockmodels (Wang and Wong, 1987; Snijders and Nowicki, 1997; Airoldi et al., 2008, 2013a) in the area of social network modeling (Goldenberg et al., 2009; Azari Soufiani and Airoldi, 2012; Choi et al., 2012). Stochastic blockmodels assume that nodes of a network are members

of latent blocks, and describe their interactions with other nodes with a parametric model. However, models for inferring modular structures from both data node variables and network data are relatively unexplored and of interest in many applications.

1.1 Contributions

In this paper, we propose an integrated probabilistic model inspired by module networks and stochastic blockmodels, to learn dependency structures from the combination of network data and node variables. We consider network data in terms of directed edges (interactions) and model network data using stochastic blockmodels. Intuitively, by incorporating complementary data types, a node which is likely to have directed edges to members of a module as well as correlation with variables of module will be assigned as parent. The use of network data enhances computational tractability and scalability of the method by restricting the space of possible dependency structures. We also show theoretically that the integration of network data leads to model identifiability, whereas node variables alone can not.

Our model captures two types of relationships between variables of modules and their parents, including small changes of variables due to global dependency structure and condition-specific large effects on variables based on parent activities in each condition. Based on these relationships, we infer a combinatorial program (Yeang and Jaakkola, 2006; Segal et al., 2003) for each module, showing how multiple parents interact in regulating the module.

For estimation of parameters, we use a Gibbs sampler instead of the deterministic algorithm employed by Segal et al. to overcome some of the problems regarding multi-modality of model likelihood (Joshi et al., 2009). We also solve the problem of sensitivity to choice of maximum number of modules using a reversible-jump MCMC method which infers the number of modules and parents based on data. The probabilistic framework infers posterior distributions of assignments of nodes to modules and thus does not face restrictions of non-overlapping modules (Airoldi et al., 2008, 2013b).

1.2 Related Work

Other works have also proposed integrating different data types, mostly as prior information, for improvement in learning structures (Werhli and Husmeier, 2007; Imoto et al., 2003; Mitra et al., 2013). Assumptions such as sparse priors have been used in other works to improve modeling of network interactions between groups of nodes (Yan et al., 2012). Our approach is different in that we consider additional data types also as observations from a model of dependency structures. Our model thus considers both network edges and node variables as data observed from the same underlying structure, providing more flexibility for the model. Moreover, we utilize data integration to identify structures between groups of nodes (modules) as opposed to individual nodes. Despite the similarity in the framework of our model to module networks, our model for variables has differences in relating modules to their parents, giving more accurate and interpretable dependencies. Also, the integration of network data is novel. Regarding the learning procedure, prior work has been done on improving module network inference by using a Gibbs sampling approach (Joshi et al., 2009). We take a step further and use a reversible-jump MCMC procedure to learn the number of modules and parents from data as well as parameter posteriors. Our method can also allow restricting the number of modules based on context, with a narrow prior. By adjusting this prior, we have multi-resolution module detection.

2. Model of Modular Structures

In the framework of module networks, dependencies are learned from profiles of node variables (e.g. gene expressions) for each node (e.g. gene), as random variables $\{X_1, ..., X_N\}$. The idea is that a group of nodes with common parents (e.g. co-regulated genes) are represented as a module and have similar probability distributions for their variables conditioned on their shared parents (regulators). Figure 1 shows a toy example where node variable data are shown in green-to-red heatmaps and network data with dashed arrows (Airoldi, 2007). A module assignment function $\mathcal{A}$ maps nodes $\{1, ..., N\}$ to K non-overlapping modules. A dependency structure function $\mathcal{S}$ assigns a set of parents Pa_j from $\{1, ..., R\}$ known candidate parents (possible regulators/influencers), which are a subset of the N nodes, to module M_j (figure 1). In the toy example, nodes d, e are assigned to the same module M_4 and b, a are assigned as their parents. In cases where multiple parents drive a module, e.g. a, b affecting M_4, combinatorial effects are represented as a decision tree (regulatory program) and each combination of parents activities, defined as a context, is assigned to a cluster of conditions (experiments). In figure 1, parent b has an activating effect while a represses M_4, hence, e, d are active in context (ii) where only b is active and a is not. Inferring this decision tree in the context of different applications shows how multiple parents act together in influencing a group of nodes, e.g. in a gene network, multiple transcription-factor (TF) proteins act as regulators together to express a group of genes.

Given this framework, our model considers variables and network data as two types of observation from the same underlying modular structure. This structure is encoded based on assignments to modules ($\mathcal{A}$) and parents for each module ($\mathcal{S}$). In the example of gene networks, in each module, TF-gene interactions are likely to be observed between TFs and upstream regions of genes in the module while combinations of expressions of TFs explain expressions of genes.

2.1 Modeling Node Variables

We model variables for nodes $\{1, ..., N\}$ in each condition or sample $c \in 1, ..., C$ with a multivariate normal represented as $\mathbf{X_c} \sim \mathcal{N}(\boldsymbol{\mu}_c, \Sigma)$, where $\mathbf{X_c}$ is a $N \times 1$ vector, with N being the total number of nodes. The covariance and mean capture two different aspects of the model regarding global dependency structures and context-specific effects of parents, respectively, as described below.

We define the covariance $\boldsymbol{\Sigma}$ to be independent of conditions and representing the strength of potential effects of one variable upon another, if the former is assigned as a parent of the module containing the latter. In the example of gene expressions, Σ may represent the affinity of a Transcription-Factor protein to a target gene promoter. The modular dependencies between variables imposes a structure on $\boldsymbol{\Sigma}$. To construct this structure, we relate node variables to their parents through a regression $\mathbf{X_c} = W\mathbf{X_c} + \boldsymbol{\epsilon}$ where $\boldsymbol{\epsilon} = \mathcal{N}(\mathbf{m_c}, I)$. W is a $N \times N$ sparse matrix in which element W_{nr} is nonzero if variable r is assigned as a parent of the module containing variable n. Here we assume W_{nr} has the same value for $\forall n \in M_k, \forall r \in Pa_k$, which leads to identifiability of model (as explained in section 6. Then, assuming $I - W$ is invertible, $\mathbf{X_c} = (I - W)^{-1}\boldsymbol{\epsilon}$ which implies $\boldsymbol{\Sigma} = (I - W)^{-T}(I - W)^{-1}$. Therefore, we impose the modular dependency structure over Σ through W, which is easier to interpret based on $\mathcal{A}, \mathcal{S}$ assignments.

We define variable means $\boldsymbol{\mu}_c$, based on parents as described below. First, based on the modular structure of nodes, we can partition the mean vector as $\boldsymbol{\mu}_c = [\boldsymbol{\mu}_c^1 ... \boldsymbol{\mu}_c^K]^T$, where each $\boldsymbol{\mu}_c^k$ for $k = 1, ..., K$ is a $1 \times N_k$ vector with N_k equal to the number of nodes in module k. In modules where there is more than one parent assigned, combinations of different activities of parents, creating a

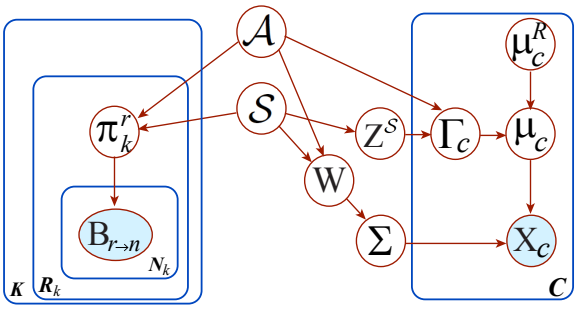

Figure 2: **Graphical representation of model:** The assignments of nodes to modules $\mathcal{A}$ and parents for modules $\mathcal{S}$ represent modular dependency structures, from which we observe node variables $\mathbf{X_c}$ in each condition c and network data $B_{r\to n}$ between a parent r and a node n. Means of node variables μ_c are determined from parent means μ_c^R with mixing coefficients Γ determined based on parent split-points Z.

context, can lead to different effects. The binary state of parent $r \in Pa_k$ is defined by comparing its mean to a split-point z_k^r, corresponding to a mixture coefficient for that state γ_{Lo}^r or γ_{Hi}^r, as: $\gamma_c^r = \gamma_{Lo}^r H(z_k^r - \mu_c^r) + \gamma_{Hi}^r H(\mu_c^r - z_k^r)$, where $H(\cdot)$ is a unit step function.

The combination of different activities are represented as a decision tree for each module k (figure 1). We represent a context-specific program as dependencies of variable means on parents activities in each context, such that $\boldsymbol{\mu}_c^k$ for module k is a linear mixture of means for parents of that module: $\boldsymbol{\mu}_c^k = \sum_{r=1}^{R_k} \boldsymbol{\gamma}_c^r \boldsymbol{\mu}_c^{Pa_k}$ where R_k is the number of parents Pa_k and $\boldsymbol{\gamma}_c^r$ are similar for all conditions c occurring in the same context. Thus, in general we can write $\boldsymbol{\mu}_c = \Gamma_c \boldsymbol{\mu}_c^R$, where $\boldsymbol{\mu}_c^R$ contains the means of parents $1, ..., R$ in condition c. The $N \times R$ matrix Γ_c has identical rows for all variables in one module based on the assignment functions $\mathcal{A}, \mathcal{S}$. The graphical model is summarized in figure 2. Thus the model for object variables would be: $\mathbf{X_c} \sim \mathcal{N}(\Gamma_c \boldsymbol{\mu}_c^R, (I - W)^{-T}(I - W)^{-1})$.

Given independent conditions, the probability of data $\mathbf{X} = [\mathbf{X_1}, ..., \mathbf{X_C}]$ for C conditions given parameters can be written as multiplication of multivariate normal distributions for each condition: $P(\mathbf{X}|\mathcal{A}, \mathcal{S}, \Theta, \Sigma, Z^S) = \prod_{c=1}^{C} P(\mathbf{X_c}|\mathcal{A}, \mathcal{S}, \theta_c, \Sigma, Z^S)$, where $\Theta = \{\theta_1, ..., \theta_C\}$ denotes the set of condition-specific parameters $\theta_c = \{\boldsymbol{\mu}_c^R, \Gamma_c\}$ for $c = 1, ..., C$ and Z^S denotes the set of parent split-points for all modules. Then for each condition we have: $P(\mathbf{X_c}|\mathcal{A}, \mathcal{S}, \theta_c, \Sigma, Z^S) = \frac{1}{(2\pi)^{N/2}|\boldsymbol{\Sigma}^{1/2}|} exp(-\frac{1}{2}(\mathbf{X_c} - \boldsymbol{\mu}_c)^T \boldsymbol{\Sigma}^{-1}(\mathbf{X_c} - \boldsymbol{\mu}_c))$.

Hence, this model provides interpretations for two types of influences of parents. By relating the distribution mean for variables in each module and in each condition to means of their assigned parents (figure 1.B), we model condition-specific effects of parents. Based on the states of parents in different contexts (partitions of conditions), this leads to a bias or large signal variations in node variables. Whereas, small signal changes (linear term) are modeled through the covariance matrix Σ which is independent of condition and is only affected by the global wiring imposed by dependency structures.

2.2 Modeling Network Data

Network data, as a directed edge between a parent $r \in \{1, ..., R\}$ and node $n \in M_k$, when r is assigned as a parent of the module $r \in Pa_k$ is defined as a directed link $B_{r \to n}$ where

$$P(B_{r \in Pa_k \to n \in M_k} | \mathcal{A}, \mathcal{S}, \pi_k^r) \sim Bernoulli(\pi_k^r) \tag{1}$$

The parameter π_k^r defines the probability of parent r influencing module M_k (figure 2). In the gene network example, an interaction between a Transcription Factor protein binding to a motif sequence, upstream of target genes, which is common in all genes of a module can be observed using ChIP data. Therefore, directed interactions from parents to all nodes in a module would be $P(B_{M_k} | \mathcal{A}, \mathcal{S}, \boldsymbol{\pi}_k) = \prod_{r \in Pa_k} \prod_{n \in M_k} P(B_{r \to n} | \mathcal{A}, \mathcal{S}, \pi_k^r)$, where $\boldsymbol{\pi_k}$ is the vector of π_k^r for all $r \in Pa_k$ and for all nodes we have:

$$\begin{aligned} P(\mathbf{B} | \mathcal{A}, S, \boldsymbol{\pi}) &= \prod_{k=1}^{K} \prod_{r \in Pa_k} \prod_{n \in M_k} P(B_{r \to n} | \mathcal{A}, \mathcal{S}, \pi_k^r) \\ &= \prod_{k=1}^{K} \prod_{r \in Pa_k} (\pi_k^r)^{s_{rk}} (1 - \pi_k^r)^{|M_k| - s_{rk}} \\ &\prod_{r' \notin Pa_k} (\pi_0)^{s_{r'k}} (1 - \pi_0)^{|M_k| - s_{r'k}} \end{aligned} \tag{2}$$

with $\boldsymbol{\pi} = \{\boldsymbol{\pi}_1, ..., \boldsymbol{\pi}_K\}$ and $s_{rk} = \sum_{n \in M_k} (B_{r \to n})$ is the sufficient statistic for the network data model and $|M_k|$ is the number of nodes in module k and π_0 is the probability that any non-parent can have interaction with a module. In gene regulatory networks, π_0 can be interpreted as basal level of physical binding that may not necessarily affect gene transcription and thus regulate a gene.

In the context of stochastic blockmodels, the group of parents assigned to each module can be considered as an individual block and thus our model can represented as overlapping blocks of nodes.

The likelihood of the model $\mathcal{M} = \{\mathcal{A}, \mathcal{S}, \Theta, \Sigma, Z^S, \boldsymbol{\pi}\}$ given the integration of node variables and network data is: $P(\mathbf{X}, \mathbf{B} | \mathcal{M}) = P(\mathbf{X} | \mathcal{A}, \mathcal{S}, \Theta, \Sigma, Z^S) P(\mathbf{B} | \mathcal{A}, \mathcal{S}, \boldsymbol{\pi})$. With priors for parameters $\mathcal{M}$ the posterior likelihood is: $P(\mathcal{M} | \mathbf{X}, \mathbf{B}) \propto P(\mathcal{M}) P(\mathbf{X}, \mathbf{B} | \mathcal{M})$.

3. Theory: Model Identifiability

Our method uses network data to avoid extra structural assumptions. In this section we formalize this idea through the identifiability of the proposed model. This property is important for interpretability of learned modules. Module networks and generally multivariate normal models for object variables can be un-identifiable, and imposing extra structural assumptions is necessary to overcome this. Here, we illustrate that the integrated learning proposed in this paper resolves the un-identifiability issue. First, we show that modeling node variables alone is identifiable only under very specific conditions. Then, we will restate some results from Latouche et al. (2011) on the identifiability of overlapping block models. Using this result we show the identifiability of the model under some reasonable conditions.

Lemma 1 Node Variables Model: *For the model of node-specific variables* $\mathbf{X}$*, if we have:* $P(\mathbf{X} | \{\mathcal{A}, \mathcal{S}\}', \Theta', \Sigma') = P(\mathbf{X} | \{\mathcal{A}, \mathcal{S}\}, \Theta, \Sigma)$

1. *Then, we can conclude:* $\mu' = \mu$ *and* $\Sigma' = \Sigma$.

2. *If we further assume* $\{\mathcal{A}, \mathcal{S}\} = \{\mathcal{A}, \mathcal{S}\}'$ *and that each module has at least two non parent nodes and* $\sum_k |Pa_k| < N$ *and the covariance matrix* Σ *is invertible, we can conclude:* $\Theta = \Theta'$, $W = W'$ *(Proof in Appendix A).*

The above lemma provides identifiability for the case where the structure $\{\mathcal{A}, \mathcal{S}\}$ is assumed to be known. However, in the case where we don't have the structure, the parameterizations of multivariate normal (μ and Σ) can be written in multiple ways in terms of Θ and $\{\mathcal{A}, \mathcal{S}\}$. This is due to existence of multiple decompositions for the covariance matrix. In the following, we will use a theorem for identifiability of overlapping block models from Latouche et al. (2011) which is an extension of the results in Allman et al. (2009). The results provide conditions for overlapping stochastic block models to be identifiable.

Theorem 1 Network Data Model: *If we have* $P(\mathbf{B}|\{\mathcal{A}, S\}, \boldsymbol{\pi}) = P(\mathbf{B}|\{\mathcal{A}, S\}', \boldsymbol{\pi}')$*, then:* $\{\mathcal{A}, S\} = \{\mathcal{A}, S\}'$ *with a permutation and* $\boldsymbol{\pi} = \boldsymbol{\pi}'$ *(except in a set of parameters which have a null Lebesgue measure) (Proof in Appendix B).*

Using the above Theorem and Lemma 1 we can have the following Theorem for the identifiability of the model.

Theorem 2 Identifiability of Model: *If we have:* $P(\mathbf{B}|\{\mathcal{A}, S\}, \boldsymbol{\pi}) = P(\mathbf{B}|\{\mathcal{A}, S\}', \boldsymbol{\pi}')$ *and* $P(\mathbf{X}|\{\mathcal{A}, \mathcal{S}\}', \Theta', \Sigma') = P(\mathbf{X}|\{\mathcal{A}, \mathcal{S}\}, \Theta, \Sigma)$ *with assuming that each module has at least two non-parent nodes and* $\sum_k |Pa_k| < N$ *and the covariance matrix* Σ *is invertible, then:* $\{\mathcal{A}, S\} = \{\mathcal{A}, S\}'$ *with a permutation,* $\boldsymbol{\pi} = \boldsymbol{\pi}'$ *,* $\Theta = \Theta'$ *and* $W = W'$ *(except in a set of parameters which have a null Lebesgue measure) (Proof in Appendix C).*

This Theorem, states the theoretical effect of integrated modeling on identifiability of modular structures, given that the sum of number of parents is less than the number of nodes (as is common in gene regulatory networks).

4. Parameter Estimation using RJMCMC

We use a Gibbs sampler to obtain the posterior distribution $P(\mathcal{M}|\mathbf{X}, \mathbf{B})$ and design Metropolis-Hastings samplers for each of the parameters $\Theta, \Sigma, \boldsymbol{\pi}$ conditioned on the other parameters and data $\mathbf{X}, \mathbf{B}$. We use Reversible-Jump MCMC (Green, 1995) for sampling from conditional distributions of the assignment and structure parameters $\mathcal{A}, \mathcal{S}$.

4.1 Learning Parameters $\Theta, \Sigma, Z^S, \boldsymbol{\pi}$.

To update the means, we only need to sample one value for means of parents assigned to the same module. This set of means of distinct parents $\boldsymbol{\mu}_c^{\mathbf{R}}$ are sampled with a normal proposal (Algorithm 1). Similarly we sample the parameters $\boldsymbol{\gamma}_c^r$, z_k^r and π_k^r, corresponding to parent $r \in Pa_k$ of module k, from normal distributions. The conditions required for identifiability (from Theorem 1) are enforced in each iteration, such that samples violating the conditions are rejected. To update covariance Σ, each distinct element of the regression matrix W corresponding to a module k, denoted as w_k, is updated. Due to the symmetric proposal distribution, the proposal is accepted with probability $P_{mh} = \min\{1, \frac{P(\mathcal{M}^{(i+1)}|X,B)}{P(\mathcal{M}^{(i)}|X,B)}\}$ where $\mathcal{M}^{(i)} = \{\mathcal{A}, \mathcal{S}, \Theta, \Sigma, Z^S\boldsymbol{\pi}\}^{(i)}$.

Algorithm 1 RJMCMC for sampling parameters

Inputs:
Node Variables Data $\mathbf{X}$
Network Data $\mathbf{B}$
for iterations $i = 1$ **to** I **do**
  Sample $\mathcal{A}^{(i+1)}$ given $\mathcal{A}^{(i)}$ using Alg 2 in appendix
  Sample $\mathcal{S}^{(i+1)}$ given $\mathcal{S}^{(i)}$ using Alg 3 in appendix
  for modules $k = 1$ **to** $K^{(i)}$ **do**
    Propose $w_k^{(i+1)} \sim \mathcal{N}(w_k^{(i)}, I)$
    Accept with probability P_{mh}; update $\Sigma^{(i+1)}$
    for parents $r = 1$ **to** R_k **do**
      Propose $z_k^{r(i+1)} \sim \mathcal{N}(z_k^{r(i)}, I)$; accept with P_{mh}
      Propose $\pi_k^{r(i+1)} \sim \mathcal{N}(\pi_k^{r(i)}, I)$; accept with P_{mh}
    end for
  end for
  for condition $c = 1$ **to** C **do**
    Propose $\boldsymbol{\mu}_c^{\mathbf{R}(i+1)} \sim \mathcal{N}(\boldsymbol{\mu}_c^{\mathbf{R}(i)}, I)$; accept with P_{mh}
    Propose $\boldsymbol{\gamma}_c^{\mathbf{R}(i+1)} \sim \mathcal{N}(\boldsymbol{\gamma}_c^{\mathbf{R}(i)}, I)$; accept with P_{mh}
  end for
end for

4.2 Learning assignments $\mathcal{A}, \mathcal{S}$.

Learning the assignment of each node to a module, involves learning the number of modules. Changing the number of modules however, changes dimensions of the parameter space and therefore, densities will not be comparable. Thus, to sample from $P(\mathcal{A}|\mathcal{S}, \Theta, \Sigma, , Z^S\boldsymbol{\pi}, \mathbf{X}, \mathbf{B})$, we use the Reversible-Jump MCMC method (Green, 1995), an extension of the Metropolis-Hastings algorithm that allows moves between models with different dimensionality. In each proposal, we consider three close move schemes of increasing or decreasing the number of modules by one, or not changing the total number. For increasing the number of modules, a random node is moved to a new module of its own and for decreasing the number, two modules are merged. In the third case, a node is randomly moved from one module to another module, to sample its assignment (Algorithm 2 in Appendix D).

To sample from the dependency structure (assignment of parents) $P(\mathcal{S}|\mathcal{A}, \Theta, \Sigma, Z^S\boldsymbol{\pi}, \mathbf{X}, \mathbf{B})$, we also implement a Reversible-Jump method, as the number of parents for each module needs to be determined. Two proposal moves are considered for $\mathcal{S}$ which include increasing or decreasing the number of parents for each module, by one (Algorithm 3 in Appendix E).

5. Results

5.1 Synthetic Data

We first tested our method on synthetic node-variables and network data generated from the proposed model. A dataset was generated for $N = 200$ nodes in $K = 4$ modules with $C = 50$ conditions for each node variable. Parents were assigned from a total of $R = 10$ number of can-

didates. Parameters π, γ and W were chosen randomly, preserving parameter sharing of modules. The inference procedure was run for 20,000 samples. Exponential prior distributions were used for number of parents assigned to each module, to avoid over-fitting. Figure 3 shows the autocorrelation

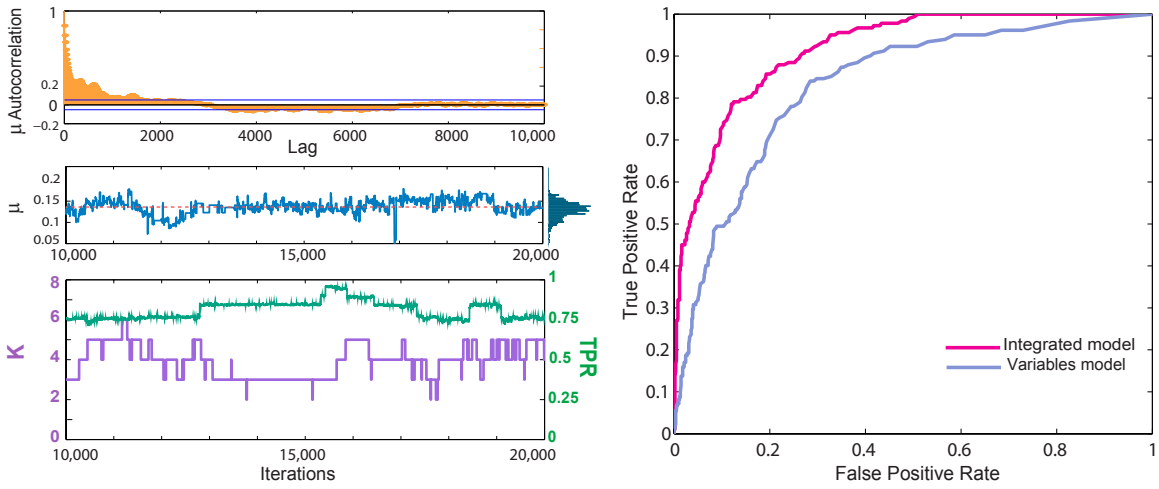

Figure 3: **Results for sythetic data:** Autocorrelation for an example variable mean (top); gibbs samples and posterior after burn-in period (actual mean shown with red line); number of modules (purple) and true positive rate of recovered links (green), ROC curve for integrated model and variables model (bottom)

for samples of variable mean μ_c^n for an example gene. The samples become independent after a lag and thus we removed the first $10,000$ iterations as burn-in period. Samples from posteriors, including the number of modules K, exhibit standard MCMC movements around the actual value (actual $K = 4$). We also calculated the true positive rate and false positive rates based on actual dependency links. We repeated the estimation of true positive and false positive rates for 100 random datasets with the same size as mentioned and computed the average ROC for the model (figure 3). As comparison, for each generated dataset, we also tested the sub-model for variable data (excluding the model for network data) to infer links (figure 3). We performed bootstrapping on sub-samples with size 1000 to compute variance of AUC (area under curve) and paired t-tests confirmed improved performance of integrated model compared to the variables sub-model ($p < 0.05$).

The parameter sharing property in modular structures allows parallel sampling of parameters w_k and $\gamma^r_{(k)}$, z^r_k,π^r_k for each module k, in each iteration and in different conditions. We used Matlab-MPI for this implementation. It takes an average of 36 ± 8 seconds to generate 100 samples for $N = 200$, $C = 50$, $R = 10$ on an i5 3.30GHz Intel(R). For further enhancement, module assignments were initialized by k-means clustering of variables.

5.2 M. tuberculosis Gene Regulatory Network

We applied our method to identify modular structures in the Mycobacterium tuberculosis (MTB) regulatory network. MTB is the causative agent of tuberculosis disease in humans and the mechanisms underlying its ability to persist inside the host are only partially known (Flynn and Chan,

2001). We used interaction data identified with ChIP-Seq of 50 MTB transcription factors and expression data for different induction levels of the same factors in 87 experiments, from a recent study by Galagan et al. (2013). Only bindings of factors to upstream intergenic regions were considered. We tested our method on 3072 MTB genes which had binding from at least one of these factors and performed 100,000 number of iterations on the combination of the two datasets. For each gene, we inferred the mode of its assignments to modules (after removing burn-in samples) and obtained 29 modules in total. The largest modules and the assigned regulators are shown in figure 4.

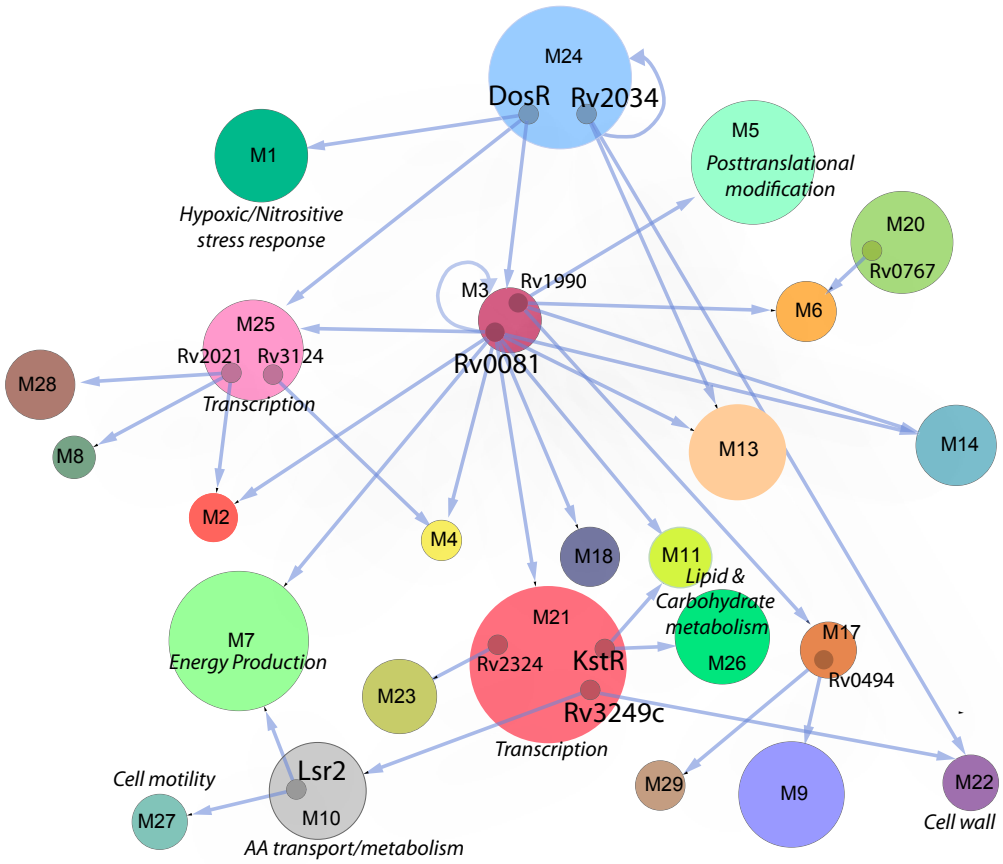

Figure 4: **Regulatory structures between largest modules inferred for MTB**: Regulators assigned to each module are shown; the size of circles are proportional to number of genes assigned to the module. Enriched functional annotations are highlighted (details in table 1).

We found functional enrichment of modules using Gene Ontology (GO) terms and COG category annotations from the TBDB database (Reddy et al., 2009) (enrichments indicate higher probability of observing a function in module compared to other modules). Out of 29 modules, 26 were enriched for at least one COG category with Bonferroni corrected $p < 0.05$. The enrichments for the top major identified modules are shown in table 1. For each module, the number of assigned genes and examples of previously studied genes are presented. The identified regulators of each

Module ID	Number of Genes	Example Genes Assigned to Module	Regulators	Enriched COG Catergories ($p < 0.05$)	Enriched GO terms ($p < 0.05$)
M21	291	KstR, Rv3249c, sigI, relA, helZ, recG	Rv0081	Replication, recombination and repair; Transcription	extracellular region; growth; plasma membrane
M24	258	DosR, sigA, sigL, clpP1,2	Rv2034	Intracellular trafficking, secretion, and vesicular transport; Secondary metabolites biosynthesis, transport and catabolism	extracellular region; plasma membrane
M7	250	Rv0324, sigE, rpoA, icl, sucC, narK1, nuoAB, nuoDEFG	Rv0081, Lsr2	Energy production and conversion; Inorganic ion transport and metabolism	NADH dehydrogenase (ubiquinone) activity; growth; plasma membrane
M5	214	inhA, fabH	Rv1990c	Posttranslational modification, protein turnover, chaperones	growth; plasma membrane
M25	161	ideR, sigB, nusG, argR, lipP, Rv2021c, Rv3124	Rv0081, DosR	Transcription; Defense mechanisms	plasma membrane; succinate dehydrogenase activity
M10	154	lysA, dapF, fprA, lipO, fadD7, fadD30, fadA6	Rv3249	Amino acid transport and metabolism;	plasma membrane
M26	148	sugA,B,C; mutA,B	KstR	Carbohydrate transport and metabolism; Lipid transport and metabolism	growth; propionate metabolic process, methylmalonyl pathway
M1	144	fabG4, fadD8	DosR	Secondary metabolites biosynthesis, transport and catabolism	cellular response to nitrosative stress; growth; plasma membrane
M22	60	fas, fadA4, pcaA, metB	Rv3249c, Rv2034	Cell wall/membrane/envelope biogenesis	plasma membrane
M27	59	kasA-B, fabD, accD6	Lsr2	Cell motility	plasma membrane
M11	48	fadA3, fadD4, lipC, lipW, nuoH-N, narI,J,H	Rv0081, KstR	Energy production and coversion; Lipid transport and metabolism	NADH dehydrogenase (ubiquinone) activity; nitrate reductase activity
M3	36	Rv0081, Rv0232, Rv1990c, fadE4, fadE5	DosR	Energy production and conversion	-

Table 1: Enrichment of functional annotations for largest modules controlled by major MTB regulators

module and enriched annotations confirm known functions for some regulators, such as the role of KstR (Rv3574) in regulating lipid metabolism (Kendall et al., 2007), confirmed in modules M26 and M11; and DosR (Rv3133c) in nitrosative stress response (Voskuil et al., 2003) (module M1) and transcription (Rustad et al., 2008) (module M25). Novel functions for other regulators and the combinations of regulators acting together are also presented.

As shown in figure 4, many modules are controlled by more than one regulator, highlighting the significance of combinatorial regulations in controlling gene expressions in this network. The inferred structure identifies multiple feed-forwards loops (FFLs), many of which involve a hub regulator Rv0081 and another regulator. FFLs are known to lead to dynamic transient responses or time delays in gene expression (Mangan and Alon, 2003) and the role of Rv0081 in driving multiple FFLs in MTB can be further studied. Also, two auto-regulating feedbacks were inferred from Rv0081 to its module M3, and from Rv2034 to M24, which may contribute to stabilizing and

noise-reduction (Kærn et al., 2005) in transcription of the hub regulators. One inferred module is M11 shown in figure 5 which is regulated by Rv0081 and KstR (Rv3574). KstR is known to be involved in cholesterol and lipid catabolism (Kendall et al., 2007) and the module is enriched for "Energy production and conversion" and "Lipid transport and metabolism" COG categories (table 1). The inferred program depicted in figure 5 shows that either of the two regulators can repress the expression of the 48 genes assigned to this module, which include lipases and genes involved in fatty acid β-oxidation and triacylglycerides cycle metabolic pathways. KstR itself is also regulated by Rv0081, forming another FFL and the roles of both factors in repressing these pathways can be further investigated. Thus, a hypoxic (oxygen deprivation) regulator Rv0081, regulates lipid metabolism genes through KstR. The two factors of hypoxic adaptation and lipid catabolism are two main factors involved in MTB persistence (Flynn and Chan, 2001; Galagan et al., 2013).

Figure 5 shows module M25 containing 161 genes, with an interesting regulatory program involving two MTB hypoxic adaptation regulators: Rv3133c (DosR) and Rv0081. DosR is well known to activate the initial response of MTB in hypoxic conditions (Park et al., 2003). As table 1 shows, M25 is enriched for "Transcription" in COG categories. The genes assigned to this module include other regulators such as Rv2021c, Rv3124 known to be induced in later time points (after 24 hours) in hypoxia. The mechanism driving this enduring hypoxic response is not well known (Rustad et al., 2008). The inferred regulatory program for this module predicts induction of most genes in the module in conditions where both DosR and Rv0081 are expressed (context (c) in figure 5). This combinatorial regulation could be acting as either a logical AND gate, where both factors are required, or Rv0081 might be the only necessary activator of the module. However, Rv0081 itself is also regulated by DosR, which creates a feed-forward loop structure driving this module (see figure 4). Hence, this program illustrates the significance of Rv0081 and DosR in the form of a FFL in mediating the induction of a second hierarchy of regulators with a time delay, leading to a later hypoxic response.

We showed in section 6 that integration of network data has theoretical advantages in terms of model identifiability. Here, we show that it can also reduce the number of false positive regulatory links in MTB data. As a gold standard, we used previously validated links (by EMSA, RTq-PCR) for two MTB regulators, including 48 known links for DosR from Voskuil et al. (2003) and 72 known links for KstR from Kendall et al. (2007). We calculated the area under precision-recall for our method by comparing posterior probabilities for DosR and KstR links to known links (table 2). As comparison, we also applied common methods shown to have best performance in DREAM challenge contests (Marbach et al., 2012) in inferring regulatory networks from gene expression only. These include Mutual Information between expression profiles (MI), CLR (Faith et al., 2007)and GENIE3 (Irrthum et al., 2010). We applied these on the above MTB expression data, and compared the inferred links to the gold standard set. As the number of validated links in MTB are small, we also scored the predictions from co-expression methods to the MTB ChIP-Seq data (Galagan et al., 2013) for the same two regulators. Also, none of these methods assume modular structures.

We then applied Module Networks (Segal et al., 2005) to the same expression dataset and compared predictions to known links and ChIP-Seq data (table 3). We set the maxmimum number of modules to 10 and constrained the candidate pool of regulators to the 50 ChIPped regulators only. On average 2.8 ± 0.63 number of regulators were assigned to each module, with a mode of 3, whereas the ChIP-Seq network shows a mode of 1 for in-degree of genes (Galagan et al., 2013), i.e. most genes have only one regulator binding. As the predicted links from module networks are deterministic, an AUPR score can not be reported, thus we compared to precision and recall of posterior

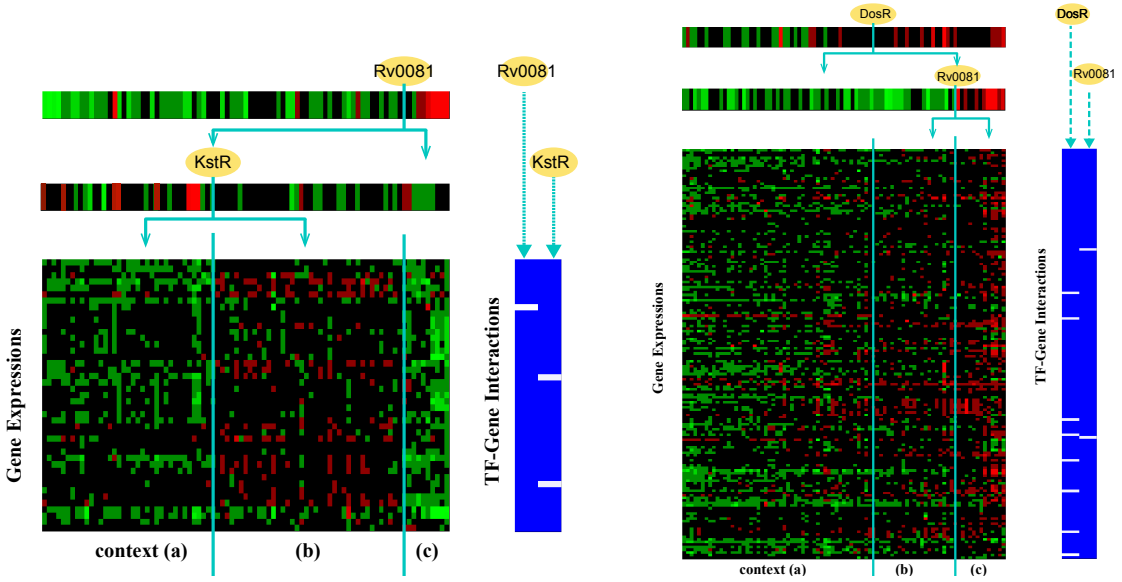

Figure 5: Examples of inferred regulatory programs: (Left) module M11 of fig. 4 showing that either of Rv0081 and KstR can repress the module in contexts (a) and (c); (Right) module M25 of fig. 4 showing the induction of these genes by DosR is mediated through Rv0081 in context (c)

Table 2: Area under precision-recall AUPR(%) calculated for link prediction using proposed method and other common co-expression methods, applied to MTB data. The predictions are scored vs known and ChIP-Seq links for two regulators

Gold Standard	Validated Links		ChIP-Seq Links	
Regulator	DosR	KstR	DosR	KstR
No. of Targets	(48)	(72)	(528)	(503)
MI	39.04	9.24	25.00	17.85
CLR	48.25	9.37	21.44	16.77
GENIE3	**62.26**	31.37	21.55	19.44
Proposed Model	**72.13**	**65.72**	**79.62**	**70.06**

mode from our models. Note small precision values are due to small number of validated links, i.e. if a link is not validated experimentally it may not be wrong. For a fair comparison of models without the effect of interaction data, we also compared to performance of our model for variables data only (table 3). These results show that module networks and in general co-expression methods have many false positives and integrating interaction data is necessary for inference of direct regulatory relationships.

Table 3: Percentage of Precision (P) and Recall (R) for link prediction using module networks and proposed models.

Gold Standard	Validated Links				ChIP-Seq Links			
Regulator	DosR		KstR		DosR		KstR	
	P	R	P	R	P	R	P	R
Module Networks	3.8	81.2	6.5	86.1	40.1	76.3	35.8	67.4
Proposed Model for Variables (mode)	4.6	77.1	7.2	77.8	55.0	83.7	52.5	80.5
Proposed Integrated Model (mode)	6.5	89.6	10.6	84.7	75.4	93.4	83.6	95.6

5.3 Twitter Network

As a second application, we used our method to find influence structures in a social network. In social networks such as twitter, the activity of users, e.g. number of tweets posted by a user in a time window, can be influenced by other users. To find these influence patterns, one approach would be to search for all other users that have correlated activity, e.g. same number of posts in a day. However, given that users are more likely to be influenced by users whom they are following, integrating the social graph of who-follows-who would improve accuracy and speed in finding influential users that affect a large community (module) of users. We applied our method on integration of two types of data from twitter. Number of user posts (tweets) are considered as node variables, time windows of one day are considered as conditions, and the network of followings is considered as network information. The dataset of tweets during a period of 4 months from (June to Sept 2009) (Yang and Leskovec, 2011) was combined with the social graph of who-follows-who (Kwak et al., 2010) and 450 number of users were randomly selected that had data in both datasets. Figure 6 shows the inferred modular structures of influence between users with circle sizes proportional to number of users assigned to each module.

The results present interesting structures of influence for each community. For example, module (community) M13 is influenced by best-selling authors such as Brian Solis and C.C. Chapman which are assigned to M5, while users in module M10 are mostly influenced by well-known web designers and developers in M3 including Ethan Marcotte (@beep), Garett Dimon and Michael Lopp (@ranks). Module M16 is mostly influenced by famous technologists including Tantek Celik (@t) and Andy Baio (@waxpancake). Module M6 contains computer scientists such as Bradley Horowitz (@elatable). The most influential users (largest fan-out degrees) in this subnetwork were famous blogger Robert Scoble (@Scobleizer) and @Starbucks. These results also identify communities with diverse interests such as M8, i.e. follow users belonging to diverse communities. The top hashtags posted by users of each community is consistent with interests and professions of their influencers and highlights major events of that period in 2009, such as launch of Google Wave, Iran election and Michael Jackson's death (figure 6). Thus this method can identify communities with common influencers in social networks.

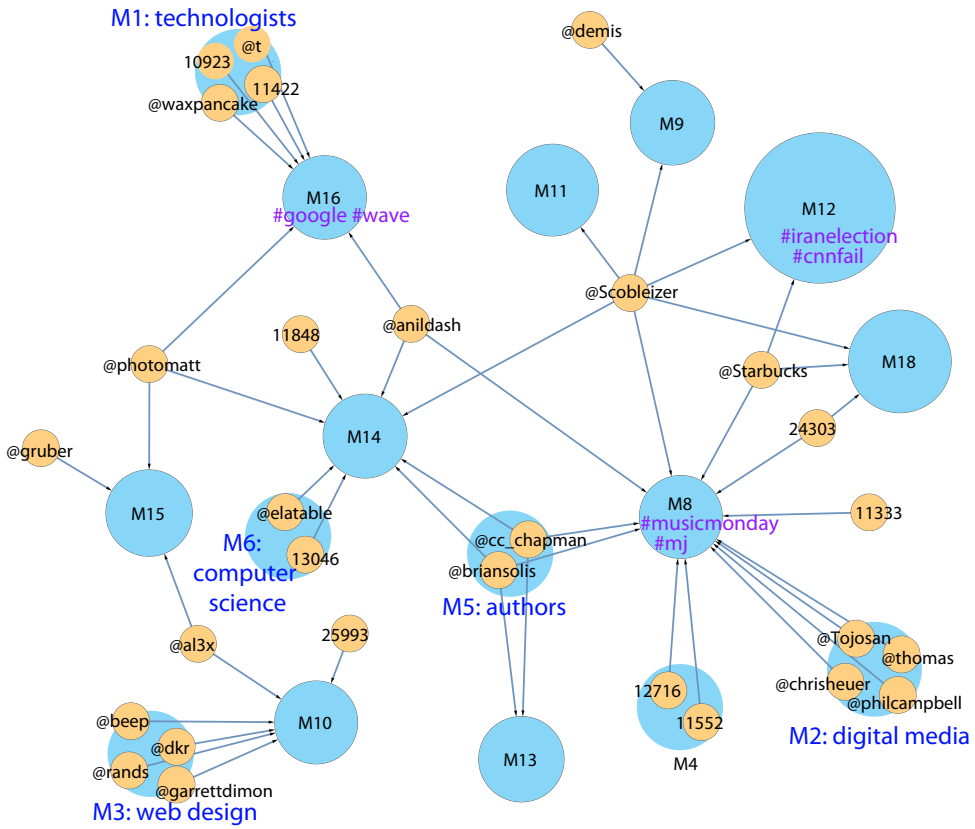

Figure 6: **Influence structures inferred for a twitter subnetwork**: Top posted hashtags and user interests are highlighted

6. Conclusion

We proposed a model for learning dependency structures between modules, from network data and node variables. We showed that the assumption of shared parents and parameters for nodes in a module, together with integration of network data deals with under-determination and un-identifiability, improves statistical robustness and avoids over-fitting. We presented a reversible-jump inference procedure for learning model posterior. Our results showed high performance on synthetic data and interpretable structures on real data from *M. tuberculosis* gene network and twitter social network. Results for MTB gene regulatory network revealed feed-forward loops and insights into condition-specific regulatory programs for lipid metabolism and hypoxic adaptation. Inferred modules in a twitter subnetwork identified influencing users for different communities.

Acknowledgments

We acknowledge funding from the Hariri Institute for Computing and Computational Science & Engineering, the National Institute of Health under grants HHSN272200800059C and R01 GM096193, the National Science Foundation under grant IIS-1149662, the Army Research Office under grant MURI W911NF-11-1-0036, and from an Alfred P. Sloan Research Fellowship.

References

E.M. Airoldi. Getting started in probabilistic graphical models. *PLoS Computational Biology*, 3 (12):e252, 2007.

E.M. Airoldi, D.M. Blei, S.E. Fienberg, and E.P. Xing. Mixed membership stochastic blockmodels. *The Journal of Machine Learning Research*, 9:1981–2014, 2008.

E.M. Airoldi, T.B. Costa, and S.H. Chan. Stochastic blockmodel approximation of a graphon: Theory and consistent estimation. In *Advances in Neural Information Processing Systems (NIPS)*, volume 26, pages 692–700, 2013a.

E.M. Airoldi, X. Wang, and X. Lin. Multi-way blockmodels for analyzing coordinated high-dimensional responses. *Annals of Applied Statistics*, 7(4):2431–2457, 2013b.

E.S. Allman, C. Matias, and J.A. Rhodes. Identifiability of parameters in latent structure models with many observed variables. *The Annals of Statistics*, 37(6A):3099–3132, 2009.

S. Aral, L. Muchnik, and A. Sundararajan. Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences*, 106(51):21544–21549, 2009.

H. Azari Soufiani and E.M. Airoldi. Graphlet decomposition of a weighted network. *Journal of Machine Learning Research*, (W&CP 22 (AISTATS)):54–63, 2012.

S.E. Celniker, L. Dillon, M.B. Gerstein, K.C. Gunsalus, S. Henikoff, G.H. Karpen, M. Kellis, E.C. Lai, J.D. Lieb, D.M. MacAlpine, et al. Unlocking the secrets of the genome. *Nature*, 459(7249): 927–930, 2009.

D.S. Choi, P.J. Wolfe, and E.M. Airoldi. Stochastic blockmodels with a growing number of classes. *Biometrika*, 99(2):273–284, Jun. 2012.

J.J. Faith, B. Hayete, J.T. Thaden, I. Mogno, J. Wierzbowski, G. Cottarel, S. Kasif, J.J. Collins, and T.S. Gardner. Large-scale mapping and validation of escherichia coli transcriptional regulation from a compendium of expression profiles. *PLoS biology*, 5(1):e8, 2007.

J.L. Flynn and J. Chan. Tuberculosis: latency and reactivation. *Infection and immunity*, 69(7): 4195–4201, 2001.

J.E. Galagan, K. Minch, M. Peterson, Anna Lyubetskaya, Elham Azizi, Linsday Sweet, Antonio Gomes, Tige Rustad, Gregory Dolganov, Irina Glotova, et al. The mycobacterium tuberculosis regulatory network and hypoxia. *Nature*, 499(7457):178–183, 2013.

A. Goldenberg, A. X. Zheng, S. E. Fienberg, and E. M. Airoldi. A survey of statistical network models. *Foundations and Trends in Machine Learning*, 2(2):129–233, Feb. 2009.

P.J. Green. Reversible jump markov chain monte carlo computation and bayesian model determination. *Biometrika*, 82(4):711–732, 1995.

S. Imoto, T. Higuchi, T. Goto, K. Tashiro, S Kuhara, , and S Miyano. Combining microarrays and biological knowledge for estimating gene networks via bayesian networks. *Proc. Computational Systems Bioinformatics*, 2003.

A. Irrthum, L. Wehenkel, P. Geurts, et al. Inferring regulatory networks from expression data using tree-based methods. *PLoS One*, 5(9):e12776, 2010.

A. Joshi, R. De Smet, K. Marchal, Y. Van de Peer, and T. Michoel. Module networks revisited: computational assessment and prioritization of model predictions. *Bioinformatics*, 25(4):490–496, 2009.

M. Kærn, T.C. Elston, W.J. Blake, and J.J. Collins. Stochasticity in gene expression: from theories to phenotypes. *Nature Reviews Genetics*, 6(6):451–464, 2005.

S.L. Kendall, M. Withers, C.N. Soffair, N.J. Moreland, S. Gurcha, B. Sidders, R. Frita, A. Ten Bokum, G.S. Besra, J.S. Lott, et al. A highly conserved transcriptional repressor controls a large regulon involved in lipid degradation in mycobacterium smegmatis and mycobacterium tuberculosis. *Molecular microbiology*, 65(3):684–699, 2007.

D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2009.

R.V. Kozinets. E-tribalized marketing?: The strategic implications of virtual communities of consumption. *European Management Journal*, 17(3):252–264, 1999.

H. Kwak, C. Lee, H. Park, and S. Moon. What is Twitter, a social network or a news media? In *WWW '10: Proceedings of the 19th international conference on World wide web*, pages 591–600, New York, NY, USA, 2010. ACM. ISBN 978-1-60558-799-8.

P. Latouche, E. Birmelé, and C. Ambroise. Overlapping stochastic block models with application to the french political blogosphere. *The Annals of Applied Statistics*, 5(1):309–336, 2011.

Y. Liu, N. Qiao, S. Zhu, M. Su, N. Sun, J. Boyd-Kirkup, and J.-D. Han. A novel bayesian network inference algorithm for integrative analysis of heterogeneous deep sequencing data. *Cell Research*, 23(3):440–443, 2013.

S. Mangan and U. Alon. Structure and function of the feed-forward loop network motif. *Proceedings of the National Academy of Sciences*, 100(21):11980–11985, 2003.

D. Marbach, J.C. Costello, R. Küffner, N.M. Vega, R.J. Prill, D.M. Camacho, K.R. Allison, M. Kellis, J.J. Collins, G. Stolovitzky, et al. Wisdom of crowds for robust gene network inference. *Nature methods*, 2012.

T. Michoel, S. Maere, E. Bonnet, Y. Joshi, A.and Saeys, T. Van den Bulcke, K. Van Leemput, P. Van Remortel, M. Kuiper, K. Marchal, et al. Validating module network learning algorithms using simulated data. *BMC Bioinformatics*, 8(Suppl 2):S5, 2007.

K. Mitra, A. Carvunis, S.K. Ramesh, and T. Ideker. Integrative approaches for finding modular structure in biological networks. *Nature Reviews Genetics*, 14(10):719–732, 2013.

H. Park, K.M. Guinn, M.I. Harrell, R. Liao, M.I. Voskuil, M. Tompa, G.K. Schoolnik, and D.R. Sherman. Rv3133c/dosr is a transcription factor that mediates the hypoxic response of mycobacterium tuberculosis. *Molecular microbiology*, 48(3):833–843, 2003.

Y. Qi and H. Ge. Modularity and dynamics of cellular networks. *PLoS Computational Biology*, 2 (12):e174, 2006.

T.B.K. Reddy, R. Riley, F. Wymore, P. Montgomery, D. DeCaprio, R. Engels, M. Gellesch, J. Hubble, D. Jen, H. Jin, et al. Tb database: an integrated platform for tuberculosis research. *Nucleic acids research*, 37(suppl 1):D499–D508, 2009.

T.R. Rustad, M.I. Harrell, R. Liao, and D.R. Sherman. The enduring hypoxic response of mycobacterium tuberculosis. *PLoS One*, 3(1):e1502, 2008.

E. Segal, M. Shapira, A. Regev, D. Pe'er, D. Botstein, D. Koller, and N. Friedman. Module networks: identifying regulatory modules and their condition-specific regulators from gene expression data. *Nature genetics*, 34(2):166–176, 2003.

E. Segal, D. Pe'er, A. Regev, D. Koller, and N. Friedman. Learning module networks. *Journal of Machine Learning Research*, (6):557–588, 2005.

T.A.B. Snijders and K. Nowicki. Estimation and prediction for stochastic blockmodels for graphs with latent block structure. *Journal of Classification*, 14(1):75–100, 1997.

M.I. Voskuil, D. Schnappinger, K.C. Visconti, M.I. Harrell, G.M. Dolganov, D.R. Sherman, and G.K. Schoolnik. Inhibition of respiration by nitric oxide induces a mycobacterium tuberculosis dormancy program. *The Journal of experimental medicine*, 198(5):705–713, 2003.

Y.J. Wang and G.Y. Wong. Stochastic blockmodels for directed graphs. *Journal of the American Statistical Association*, 82(397):8–19, 1987.

A.V. Werhli and D. Husmeier. Reconstructing gene regulatory networks with bayesian networks by combining expression data with multiple sources of prior knowledge. *Statistical Applications in Genetics and Molecular Biology*, 6(1):15, 2007.

F. Yan, Z. Xu, and Y. Qi. Sparse matrix-variate gaussian process blockmodels for network modeling. *arXiv preprint arXiv:1202.3769*, 2012.

J. Yang and J. Leskovec. Patterns of temporal variation in online media. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 177–186. ACM, 2011.

C.-H. Yeang and T. Jaakkola. Modeling the combinatorial functions of multiple transcription factors. *Journal of Computational Biology*, 13(2):463–480, 2006.

C.-H. Yeang, T. Ideker, and T. Jaakkola. Physical network models. *Journal of computational biology*, 11(2-3):243–262, 2004.

Appendix A. Proof of Lemma 1

Lemma 1 Node Variables Model: *For the model of node variables* $\mathbf{X}$*, if we have:*

$$P(\mathbf{X}|\{\mathcal{A},\mathcal{S}\}',\Theta',\Sigma') = P(\mathbf{X}|\{\mathcal{A},\mathcal{S}\},\Theta,\Sigma) \tag{3}$$

1. *Then, we can conclude:* $\mu' = \mu$ *and* $\Sigma' = \Sigma$*.*
2. *If we further assume* $\{\mathcal{A},\mathcal{S}\} = \{\mathcal{A},\mathcal{S}\}'$ *and that each module has at least two non parent nodes and* $\sum_k |Pa_k| < N$ *and the covariance matrix* Σ *is invertible, we can conclude:* $\Theta = \Theta'$*,* $W = W'$*.*

Proof sketch:

1. Considering that distributions of $\mathbf{X}$ are multivariate Normal under both parameter sets, it is straight forward that the mean and covariance parameters of two Normals should be the same. This can be formally shown by finding maximum of the distribution and curvature at any point for both sides, hence, $\mu' = \mu$ and $\Sigma' = \Sigma$.

2. From the identifiability of μ and Σ, it is sufficient to show that μ and Σ uniquely define Θ, W given $\{\mathcal{A},\mathcal{S}\}$. Starting from Γ_c, we can consider the following set of linear equations:

   $$\boldsymbol{\mu}_c = \Gamma_c \boldsymbol{\mu}_c^R$$

   This is a set of equations with N equations and $\sum_k |Pa_k|$ unknowns. Hence, when $\sum_k |Pa_k| < N$ this set of linear equations will lead to a unique solution if a solution exists.

   For the Σ, given that it is invertible, we have:

   $$\Sigma^{-1} = (I - W)^T (I - W) \tag{4}$$

   Considering that parents have the same value for W_{nr} for $\forall n \in M_k$. Then, we can simply find W_{nr} by solving $|Pa_k| * {W_{nr}}^2 = \Sigma^{-1}_{ij}$ where i, j are two genes that are non parents and belong to the module M_k.

□

Appendix B. Proof of Theorem 1

Theorem 1 Network data Model: *If we have:*

$$P(\mathbf{B}|\{\mathcal{A},S\},\boldsymbol{\pi}) = P(\mathbf{B}|\{\mathcal{A},S\}',\boldsymbol{\pi}') \tag{5}$$

Then: $\{\mathcal{A},S\} = \{\mathcal{A},S\}'$ *with a permutation and* $\boldsymbol{\pi} = \boldsymbol{\pi}'$*(except in a set of parameters which have a null Lebesgue measure).*

Proof sketch: Our network data model is an overlapping stochastic block model, where the blocks are parents and modules, with a specific parametrization among the modules and parents. Hence, we have the identifiability using the Theorem 4.1 in Latouche et al. (2011). □

Appendix C. Proof of Theorem 2

Theorem 2 Identifiability of model: *If we have:*

$$P(\mathbf{B}|\{\mathcal{A}, S\}, \boldsymbol{\pi}) = P(\mathbf{B}|\{\mathcal{A}, S\}', \boldsymbol{\pi}') \tag{6}$$
$$P(\mathbf{X}|\{\mathcal{A}, \mathcal{S}\}', \Theta', \Sigma') = P(\mathbf{X}|\{\mathcal{A}, \mathcal{S}\}, \Theta, \Sigma) \tag{7}$$

with assuming that each module has at least two non-parent nodes and $\sum_k |Pa_k| < N$ *and the covariance matrix* Σ *is invertible, then:* $\{\mathcal{A}, S\} = \{\mathcal{A}, S\}'$ *with a permutation,* $\boldsymbol{\pi} = \boldsymbol{\pi}'$ *,* $\Theta = \Theta'$ *and* $W = W'$*.*

Proof sketch: This theorem is an immediate result from combination of Theorem 1 and Lemma 1. Using (6), according to Theorem 1 we have: $\{\mathcal{A}, S\} = \{\mathcal{A}, S\}'$ with a permutation and $\boldsymbol{\pi} = \boldsymbol{\pi}'$. Now, knowing $\{\mathcal{A}, S\} = \{\mathcal{A}, S\}'$ and equation (7) we can apply Lemma 1 leading to $\Theta = \Theta'$ and $W = W'$. This concludes the proof. □

Appendix D. Learning Module Assignment $\mathcal{A}$.

Learning the assignment of each gene to a module, involves learning the number of modules. Changing the number of modules however, changes dimensions of the parameter space and therefore, densities will not be comparable. Thus, to sample from $P(\mathcal{A}|\mathcal{S}, \Theta, \Sigma, , Z^S\boldsymbol{\pi}, \mathbf{X}, \mathbf{B})$, we use the Reversible-Jump MCMC method (Green, 1995), an extension of the Metropolis-Hastings algorithm that allows moves between models with different dimensionality.

In each proposal, we consider three close move schemes of increasing or decreasing the number of modules by one, or not changing the total number. For increasing the number of modules, a random gene is moved to a new module of its own and for decreasing the number, two modules are merged. In the third case, an gene is randomly moved from one module to another module, to sample its assignment.

We design transformation of parameters using Green's method to extend model dimensions (Algorithm 2) The acceptance ratio for the split move is $P_{split} = \min\{1, \frac{P(\mathcal{M}^{(i+1)}|X,B)}{P(\mathcal{M}^{(i)}|X,B)} \times \frac{\frac{1}{K+1}}{\frac{1}{K}} \times \frac{p_{+1}}{p_{-1}} \times \frac{1}{p(\mathbf{u})p(\mathbf{u}')} \times \mathcal{J}_{(i)\to(i+1)}\}$ where $\mathcal{J}_{(i)\to(i+1)}$ is the Jacobian of the transformation from the previous state to the proposed state, and the acceptance ratio for the merge move is $P_{merge} = \min\{1, \frac{P(\mathcal{M}^{(i+1)}|X,B)}{P(\mathcal{M}^{(i)}|X,B)} \times \frac{\frac{1}{K-1}}{\frac{1}{K}} \times \frac{p_{-1}}{p_{+1}} \times \mathcal{J}_{(i)\to(i+1)}\}$.

Algorithm 2 RJMCMC to update $\mathcal{A}$

Find K: number of distinct modules in $\mathcal{A}^{(i)}$
Propose move ν from $\{-1, 0, +1\}$ with probabilities p_{-1}, p_0, p_{+1}, respectively.
switch (ν)
case +1**:**
  Select random gene $n \in M_k$ uniformly
  Assign n to new module M_{K+1}
  Assign parents $Pa_{K+1} = Pa_k$
  Draw vectors $\mathbf{u}, \mathbf{u}' \sim \mathcal{N}(0, 1)$
  Propose parameters:
  $\boldsymbol{\pi}_{k1}^{Pa_{K+1}} = \boldsymbol{\pi}_k^{Pa_k} - \mathbf{u}$, $\boldsymbol{\pi}_{k2}^{Pa_k} = \boldsymbol{\pi}_k^{Pa_k} + \mathbf{u}$
  $\boldsymbol{\gamma}_{k1}^{Pa_{K+1}} = \boldsymbol{\gamma}_k^{Pa_k} - \mathbf{u}'$, $\boldsymbol{\gamma}_{k2}^{Pa_k} = \boldsymbol{\gamma}_k^{Pa_k} + \mathbf{u}'$
  Compute $\{\Theta, \Sigma, \boldsymbol{\pi}\}$
  Accept $\mathcal{A}^{(i+1)}$ with P_{split}
case -1**:**
  Select two random modules M_{k_1} and M_{k_2}
  Merge into one module $M_k 1$
  Assign parents $Pa_{k1} = Pa_{k_1} \cup Pa_{k_2}$
  for $\forall r \in Pa_{k_1} \cap Pa_{k_2}$ **do**
    Propose $\boldsymbol{\pi}_{k1}^r = (\boldsymbol{\pi}_{k1}^r + \boldsymbol{\pi}_{k2}^r)/2$
    and $\boldsymbol{\gamma}_{k1}^r = (\boldsymbol{\gamma}_{k1}^r + \boldsymbol{\gamma}_{k2}^r)/2$
  end for
  Compute $\{\Theta, \Sigma, \boldsymbol{\pi}\}$
  Accept $\mathcal{A}^{(i+1)}$ with P_{merge}
case 0**:**
  Select two random modules M_{k_1}, M_{k_2}
  Move a random gene n from M_{k_1} to M_{k_2}
  Compute $\{\Theta, \Sigma, \boldsymbol{\pi}\}$
  Accept $\mathcal{A}^{(i+1)}(n) = k_2$ with P_{mh}
end switch

Appendix E. Learning Dependency Structure $\mathcal{S}$.

To sample from the dependency structure $P(\mathcal{S}|\mathcal{A}, \Theta, \Sigma, Z^S\boldsymbol{\pi}, \mathbf{X}, \mathbf{B})$ (assignment of parents), we also implement a Reversible-Jump method, as the number of parents for each module needs to be determined. Two proposal moves are considered for $\mathcal{S}$ which include increasing or decreasing the number of parents for each module, by one (Algorithm 3). In the case of addition of a parent to a module, we propose mixture coefficients γ and interaction parameters π for the added regulator, based on its learned values in another module, where it has already been assigned as a parent, with an additional noise term. The acceptance ratio for the add proposal is $P_{add} = \min\{1, \frac{P(\mathcal{M}^{(i+1)}|X,B)}{P(\mathcal{M}^{(i)}|X,B)} \times \frac{\frac{1}{R_k+1}}{\frac{1}{R-R_k}} \times \frac{p_S}{1-p_S} \times \frac{1}{p(u)p(\mathbf{u}')} \times \mathcal{J}_{(i)\to(i+1)}\}$ where R_k is the number of parents for module k in the

$i-$th state, and the acceptance ratio for the remove proposal is $P_{rem} = \min\{1, \frac{P(\mathcal{M}^{(i+1)}|X,B)}{P(\mathcal{M}^{(i)}|X,B)} \times \frac{\frac{1}{R-R_k+1}}{\frac{1}{R_k}} \times \frac{1-p_S}{p_S} \times \mathcal{J}_{(i)\to(i+1)}\}$.

Algorithm 3 RJMCMC to update $\mathcal{S}$

Set p_S
for module $k = 1$ **to** K **do**
  Propose ν from $\{+1, -1\}$ with p_S
  switch (ν)
  case $+1$**:**
    Add a random parent $r \in 1, ..., R$ to Pa_k
    Draw $u, \mathbf{u}' \sim Unif(0, 1)$
    if r is also a parent of another module $Pa_{k'}$ **then**
      Propose $\pi_k^r = \pi_{k'}^r + u$, $\gamma_c^{r_k} = \gamma_c^{r_{k'}} + \mathbf{u}'(c)$ for all $c \in \{1, ..., C\}$
    else
      Propose $\pi_k^r = u$,$\gamma_c^{r_k} = \mathbf{u}'(c)$ for all c
    end if
    Compute $\{\Theta, \Sigma, \boldsymbol{\pi}\}$
    Accept $\mathcal{S}^{(i+1)}$ with P_{add}
  case -1**:**
    Remove a random parent r from Pa_k
    Compute $\{\Theta, \Sigma, \boldsymbol{\pi}\}$
    Accept $\mathcal{S}^{(i+1)}$ with P_{rem}
  end switch
end for